arXiv:1701.05863v1 [stat.AP] 20 Jan 2017

Statistical Analysis of Origin-Destination Point Patterns: Modeling Car Thefts and Recoveries

Shinichiro Shirota*, Jorge Mateu† and Alan. E. Gelfand‡

January 23, 2017

Abstract

Origin-destination problems have a substantial literature at this point. Applications have included labor flows, traffic flows, and mail flows. Viewed from a spatial perspective, the data consist of a pair of spatial locations, the origin and the destination with an associated measurement. Customarily, the locations are viewed as areal units. However, here the data we seek to explain consist of a point pattern of origins and a partial point pattern of associated destinations. The point pattern of origins is a collection of car theft locations. The point pattern of destinations is the partial collection of recovery locations.

We offer four modeling approaches for analyzing such data. We apply them to two datasets. One is small from the state of Neza in Mexico but with areal covariate information regarding population features and crime type. A second, much larger one, is from Belo Horizonte in Brazil but lacks covariates.

Keywords: log Gaussian Cox process, nonhomogeneous Poisson process, posterior predictive distribution, rank probability score, trend surface

1 Introduction

Mobility and spatial interaction data have become increasingly available due to the wide adoption of location-aware technologies (Guo et al. (2012)). Examples of mobility data include human daily activities, vehicle trajectories, and animal movements. A special type of mobility data can be found in origin-destination problems which have a substantial literature at this point (Wood et al. (2010); Adrienko and Adrienko

*Department of Statistical Science, Duke University, US. E-mail: ss571@stat.duke.edu
†Department of Mathematics, Universitat Jaume I, Spain. E-mail: mateu@mat.uji.es
‡Department of Statistical Science, Duke University, US. E-mail: alan@stat.duke.edu

(2011); Guo et al. (2012)). Applications noted below have included labor flows, traffic flows, and mail flows. Customarily, such data is collected at areal unit scale with *distance* expected to influence flow. An example might be regular travel paths based on commuting between residence and workplace. Vehicle trajectories, e.g., taxi rides might give rise to point-based trajectories that only have the origin and destination points for each passenger ride. Our application here consists of origin-destination points providing locations of car thefts and car recoveries with, of course, no knowledge on the trajectory in between (Assunção and Lopes (2007); Lopes and Assunção (2012)). Homicide is a more complex linked events situation since we can have at least three linked positions: the murder location, and the victim and offender addresses (see details in Lopes and Assunção (2012)). What each of these situations have in common is that two or more spatial point patterns are linked by an overall event.

In this article we focus on pointed-referenced origin-destination pairs. Viewed from a spatial perspective, the data consist of a pair of spatial locations, the origin and the destination say with a measurement at each location and a distance between them. Typically, locations are areal units, e.g., postal codes where the measurements might be volumes in each post code along with delivery times from one to another (Banerjee et al. (2000)) or, for labor markets, the measurements might be the number of individuals living in one unit, working in another with commuting time between them (Chakraborty et al. (2013)).

We focus here on the origin-destination problem above where the data consists of a point pattern of origins and a partial point pattern of associated destinations. In particular, the point pattern of origins is a collection of car theft locations. The point pattern of destinations is a collection of recovery locations. In practice, we have many fewer origin-destination pairs than origin locations; that is, only a fraction of the stolen vehicles are recovered.

We are motivated by two real data settings. One consists of a collection of automobile thefts, with a fraction of recoveries, over the state of Neza in Mexico. We have a total of 4,016 car theft locations (after deleting some missing locations) during 2015, over both northern and southern parts of Neza. This dataset is small but is endowed with areal covariate information regarding population features and crime type that can be used for explanation in our modeling strategy. See Figures 1 and 2, and Section 2 for further description. A second dataset consists of car thefts which occurred in Belo Horizonte (Brazil). It is a much larger dataset, but lacks covariates. This city is 331 km^2 in area and has approximately 2.4 million inhabitants. In the period from August, 1, 2000, to July, 31, 2001, the dataset consists of 5,250 pairs of theft and recovery locations. See Figure 3 and Section 2 for a more complete description.

The connection between origin-destination problems and spatial point processes has hardly been treated in the literature. Beneš et al. (2005) consider statistical

analysis of linked point processes, where for each case of a disease in their study, they have the coordinates of the individual's home and of the reported infection location. However, they used only the distance between the two linked locations. As noted above, Banerjee et al. (2000) analyzed postal service performance with regard to delivery of priority mail. However, their data structure was aggregated to Swiss postal regions, not of the point process type. Assunção and Lopes (2007) and Lopes and Assunção (2012) consider bivariate linked point processes as point processes with events marked with another spatial event representing origin-destination data types. Their methods are illustrated with the Belo Horizonte data on car theft locations and the eventual car retrieval locations, also analyzed here.

We offer four statistical modeling approaches to analyze such origin-destination data. First, we model the set of origin locations using both a nonhomogeneous Poisson process as well as a log-Gaussian Cox process. We demonstrate the benefit of the latter specification. Second, we consider a conditional specification to provide the distribution of destination locations given origin locations. Third, we consider a joint model, viewing the data as an origin-destination pair, and treating the point pattern as consisting of random pairs of locations. Because both origin and destination are points in $\mathbb{R}^2$, we propose to specify the model as a point pattern over a bounded set $D_o \times D_d \subset \mathbb{R}^2 \times \mathbb{R}^2$. This approach needs an intensity over $D_o \times D_d$ linking pairs of locations. In this regard, the destination can be viewed as a mark associated with the origin where the mark lies in a subset of $\mathbb{R}^2$. Finally, in our fourth approach, we reduce this joint origin-destination model to a specification over a subset of $\mathbb{R}^1 \times \mathbb{R}^1$ where, say for theft locations, different from the first approach, we model the x- and y- coordinates of a theft location using special cases of trend surface models which yields intensities with *ridges*, enabling potentially useful interpretations. All these statistical approaches provide a deeper look at the analysis of the motivating examples.

The plan of the paper is the following. Section 2 provides a description of the two datasets that motivate this paper. Section 3 presents the statistical approach that models the set of origin locations using both a nonhomogeneous Poisson process as well as a log-Gaussian Cox process. Then Section 4 considers the conditional specification approach that provides the distribution of destination locations given origin locations. Section 5 supplies the joint model approach, viewing the data as an origin-destination pair, and treating the point pattern as consisting of random pairs of locations. Finally, Section 6 presents the analogous ridge (trend surface) model specification. The paper ends with a summary and future work.

2 Data Description

This paper presents the analysis of two real datasets consisting of a collection of automobile thefts and recoveries over the state of Neza in Mexico, and also car thefts occurred in Belo Horizonte (Brazil).

2.1 The Neza data

The Ciudad Neza (referred to as Neza in what follows) is a city and municipality adjacent to the northeast corner of Mexico's Federal district. It is part of the Mexico City metropolitan area. The region is composed of the north and south parts separated by a single road. On the east side of this road there is a large park, and on the west side an airport. In the analysis below, we separate these two regions. Our dataset contains car theft locations in 2015. The number of car theft locations is 4,016 after deleting some missing locations.

We also have several areal unit covariates split into two categories. The first category consists of population types: (1) `Pop15` - number of individuals 15 years and older, (2) `Apart` - number of apartments, (3) `Eco` - number of economically active individuals, (4) `Employ` - number of employed individuals, hence `unEmploy` - number of unemployed individuals, (5) `inBorn` - number of individuals born in the area, hence `outBorn` - number of individuals born outside the area, (6) `Health` - number of individuals with health insurance access, hence `noHealth` - number of individuals without health insurance access and (7) `Scholar` average of scholarly grade (integer values from $6-10$). The second category consists of crime types: (1) `Extor` - number of extortion crimes, (2) `Murder` - number of murders, (3) `Burg` - number of burglaries, (4) `Shop` - number of shop robberies, (5) `Public` - number of public transport robberies, (6) `Street` - number of street robberies, (7) `Kidnap` - number of kidnappings and (8) `Total` - total number of infractions (some additional crimes beyond (1) through (7) are included here). These covariates are provided for 90 disjoint blocks in Neza. Figure 1 shows the theft locations for the north and south regions. The x-axis and y-axis are in eastings and northings. 22 blocks are located in the north region with the remaining 68 blocks in the south region. They are indicated in white in the figure. Of the thefts, 3,327 points (689 points) are observed in the south (north) region. These locations seem to be spread smoothly over the each region rather than suggesting concentration in "hot spots."

Unfortunately, the recovery locations are observed for only 382 theft locations. 56 of these locations are outside the Neza region. Figure 2 shows the plot of the recovery locations for observed theft and recovery pairs as well as a histogram of the distance between theft and recovery locations. We see that the recovery location tends to be near the theft location. This motivates our ensuing modeling.

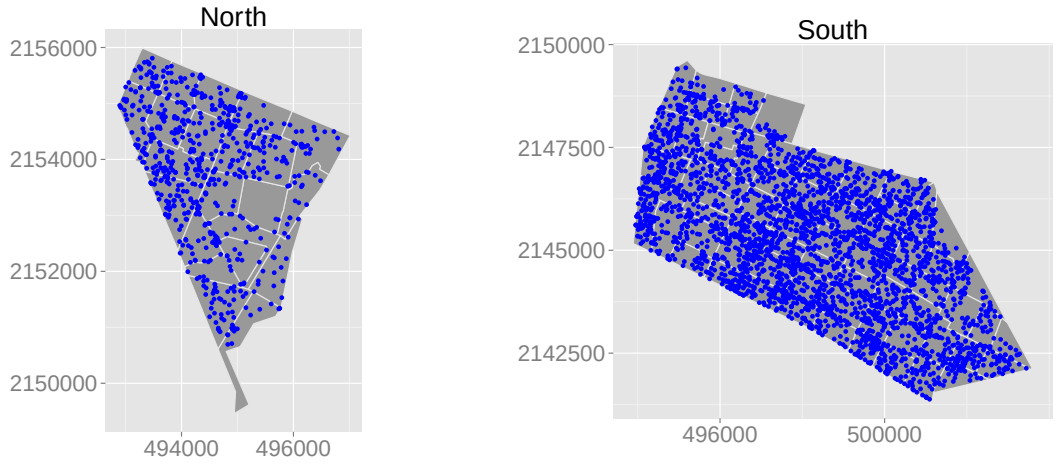

Figure 1: Car theft locations in the north region (left) and the south region (right) in Neza

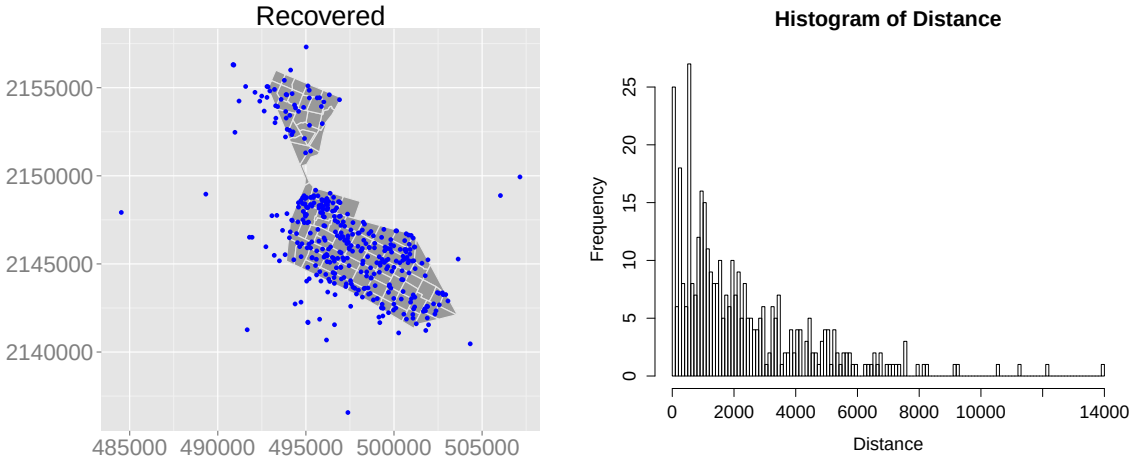

Figure 2: Recovery locations (left) and histogram of the distance between theft and recovery locations (right)

2.2 The Belo Horizonte data

We also examine car theft and recovery point patterns in Belo Horizonte in Brazil (Lopes and Assunção (2012)). This dataset contains 5,250 pairs of theft and recovery locations, an order of magnitude larger than we have in the Neza data. However, this dataset does not have any covariate information. The left panel of Figure 3 shows the point patterns of theft and recovery locations. The point patterns are similar, though recovery points seemed to be a bit more concentrated. The right panel provides the histogram of the distance between theft and recovery locations. Again, recovery location tends to be near theft location; in fact, 770 pairs (roughly 15%) are observed to be within 200 m.

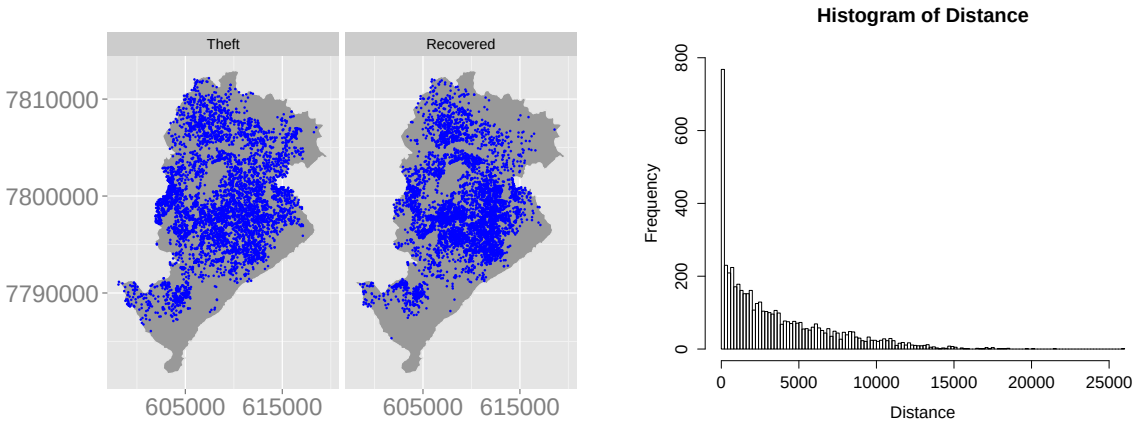

Figure 3: Car theft and recovery locations (left) and histogram of the distance between theft and recovery locations (right) in Belo Horizonte

3 Modeling of car thefts

3.1 LGCP and NHPP models for vehicle theft

As our first analysis, we consider the vehicle theft events in Neza. This is because we have a relatively large number of theft points in Neza and also have covariate information. Let $\mathcal{S} = \{\boldsymbol{s}_1, \ldots, \boldsymbol{s}_n\}$ denote the observed point pattern over the study region $D \subset \mathbb{R}^2$. In our case, $\mathcal{S}$ is the set of car theft locations and D is the north or the south region. We consider non-homogeneous Poisson processes (NHPP) and log Gaussian Cox processes (LGCP, Møller et al. (1998)) for modeling theft events. As a spatial process model, the LGCP is defined so that the log of the intensity is a

Gaussian process (GP), i.e.,

$$\log \lambda(\boldsymbol{s}) = \boldsymbol{X}(\boldsymbol{s})\boldsymbol{\beta} + z(\boldsymbol{s}), \quad \boldsymbol{z}(\mathcal{S}) \sim \mathcal{N}(\boldsymbol{0}, \mathbf{C}_{\boldsymbol{z}}), \quad \boldsymbol{s} \in D. \tag{1}$$

where $\boldsymbol{X}(\boldsymbol{s})$ is a covariate vector at $\boldsymbol{s}$ and $z(\boldsymbol{s})$ is a Gaussian process. In particular, $\boldsymbol{z}(\mathcal{S}) = (z(\boldsymbol{s}_1), \ldots, z(\boldsymbol{s}_n))$ denotes the n-variate zero mean Gaussian distribution, with covariance matrix $\mathbf{C}_{\boldsymbol{z}} = (C(\boldsymbol{s}_i, \boldsymbol{s}_j))_{i,j=1,\ldots,n}$, which provides spatial random effects for the intensity surface, pushing up and pulling down the surface, as appropriate. We assume an exponential covariance function, i.e., $C(\boldsymbol{u}, \boldsymbol{u}') = \sigma^2 \exp(-\phi \|\boldsymbol{u} - \boldsymbol{u}'\|)$[1].

If we remove $\boldsymbol{z}(\boldsymbol{s})$ from the log intensity, we obtain the associated NHPP. NHPP's have a long history in the literature (see, e.g., Illian et al. (2008)). Furthermore, given $\lambda(\boldsymbol{s})$ with $z(\boldsymbol{s})$, $\mathcal{S}$ is, again, an NHPP with intensity $\lambda(\boldsymbol{s})$. The likelihood takes the form

$$\mathcal{L}(\mathcal{S}) \propto \exp\left(-\int_D \lambda(\boldsymbol{u}) d\boldsymbol{u}\right) \prod_{i=1}^{n} \lambda(\boldsymbol{s}_i) \tag{2}$$

For inference with a LGCP using (2), we need to approximate the stochastic integral inside the exponential. We create K grid cells roughly uniformly over the study region D; convergence to the exact posterior distribution when $K \to \infty$ is guaranteed following Waagepetersen (2004). Then, the approximate likelihood for the LGCP becomes

$$\mathcal{L}(\mathcal{S}) \propto \exp\left(-\sum_{k=1}^{K} \lambda(\boldsymbol{u}_k)\Delta_k\right) \prod_{k=1}^{K} \lambda(\boldsymbol{u}_k)^{n_k} \tag{3}$$

where n_k is the number of points in k-th grid, i.e., $\sum_k^K n_k = n$, Δ_k is the area of k-th grid (in practice, we standardize Δ_k so that $\sum_k^K \Delta_k = |D| = 1$) and u_k is the "representative point" for k-th grid (e.g., Banerjee et al. (2014) and Møller and Waagepetersen (2004)). In fact, since we have covariate values for 90 different areal units, we adopt this discretization. In order to implement full inference we adopt a Bayesian framework, fitting the model using Markov chain Monte Carlo (see, e.g., Robert and Casella (2004)), as discussed below.

3.2 Covariate selection

The spatstat R-package (Baddeley and Turner (2005); Baddeley et al. (2013)) supports the model fitting of spatial point processes, in particular Poisson processes, and related inference and diagnostic tools. The function `ppm` fits a spatial point process

[1]Since we never see observations of the intensity, it is hard to justify or identify a richer covariance function for the spatial random effects.

to an observed point pattern and allows the inclusion of covariates. Two estimation methods, pseudo-maximum likelihood (Baddeley and Turner (2000)) and approximate maximum likelihood (Huang and Ogata (1999)), are implemented. Working with the 90 (22 in north and 68 in south) blocks, we select covariates from those listed in Section 2.1 by forward and backward selection (`step` function in R) based on the models fitted by the `ppm` function. The resulting selection is presented in Section 3.4.

3.3 p-thinning cross validation

Cross validation is a standard approach for assessing model adequacy and is available for point pattern models with conditionally independent locations given the intensity, hence, for both the NHPP and LGCP (see, Leininger and Gelfand (2016)).

We implement cross validation by obtaining a training (fitting) dataset and a testing (validation) dataset using p-thinning as proposed by Leininger and Gelfand (2016). Let p denote the retention probability, i.e., we delete $\boldsymbol{s}_i \in \mathcal{S}$ with probability $1-p$. This produces a training point pattern $\mathcal{S}^{train}$ and a test point pattern $\mathcal{S}^{test}$, which are independent, conditional on $\lambda(\boldsymbol{s})$. In particular, $\mathcal{S}^{train}$ has intensity $\lambda(\boldsymbol{s})^{train} = p\lambda(\boldsymbol{s})$. We set $p = 0.5$ and estimate $\lambda(\boldsymbol{s})^{train}$ $\boldsymbol{s} \in D$. Then, we convert the posterior draws of $\lambda^{train}(\boldsymbol{s})$ into predictive draws of $\lambda^{test}(\boldsymbol{s})$ using $\lambda^{test}(\boldsymbol{s}) = \frac{1-p}{p}\lambda^{train}(\boldsymbol{s}) = \lambda^{train}(\boldsymbol{s})$.

Let $\{B_r\}$ be a collection of subsets of D as a evaluation grid. For the choice of $\{B_r\}$, Leininger and Gelfand (2016) suggest to draw random subsets of the same size uniformly over D. Specifically, for $q \in (0,1)$, if the area of each B_r is $q|D|$, then q is the *relative* size of each B_r. They argue that making the subsets disjoint is time consuming and unnecessary. Based on the p-thinning cross validation, we consider two model performance criteria: (i) predictive interval coverage (PIC) and (ii) rank probability score (RPS). PIC offers assessment of model adequacy, RPS enables model comparison.

Predictive Interval Coverage

After the model is fitted to $\mathcal{S}^{train}$, the posterior predictive intensity function can supply posterior predictive point patterns and therefore samples from the posterior predictive distribution of $N(B_r)$ for each r. For the ℓ-th posterior sample, $\ell = 1, ..., L$, the associated predictive residual is defined as

$$R_\ell^{pred}(B_r) = N^{test}(B_r) - N^{(\ell)}(B_r) \tag{4}$$

where $N^{test}(B_r)$ is the number of points of the test data in B_r. If the model is adequate, the empirical predictive interval coverage rate, i.e., the proportion of intervals

which contain 0, is expected to be roughly the nominal level of coverage; below, we choose 90% nominal coverage. Empirical coverage much less than the nominal suggests model inadequacy; predictive intervals are too optimistic. Empirical coverage much above, for example 100%, is also undesirable. It suggests that the model is introducing more uncertainty than needed.

Rank Probability Score

Gneiting and Raftery (2007) propose the continuous rank probability score (CRPS). This score is derived as a proper scoring rule and enables a criterion for assessing the precision of a predictive distribution for continuous variables. In our context, we seek to compare a predictive distribution to an observed count. Czado et al. (2009) and refeences therein discuss rank probability scores (RPS) for count data. Intuitively, a good model will provide a predictive distribution that is very concentrated around the observed count. While the RPS has a challenging formal computational form, it is directly amenable to Monte Carlo integration. In particular, for a given B_r, we calculate the RPS as

$$\begin{aligned} \text{RPS}(F, N^{test}(B_r)) = & \frac{1}{L}\sum_{\ell=1}^{L} |N^{(\ell)}(B_r) - N^{test}(B_r)| \\ & - \frac{1}{2L^2}\sum_{\ell=1}^{L}\sum_{\ell'=1}^{L} |N^{(\ell)}(B_r) - N^{(\ell')}(B_r)| \end{aligned} \tag{5}$$

Summing over the collection of B_r gives a model comparison criterion. Smaller values of the sum are preferred.

3.4 Results for theft locations

We implement model fitting with the NHPP and the LGCP for theft locations in the two separate regions in Neza. Again, $K_N = 22$ and $K_S = 68$ are the number of grid cells for the north and south regions, respectively. We rescale the northings and eastings dividing by 1,000. The estimation results are presented at this scale.

Working with the NHPP model, we implement the forward and backward algorithm (`step` function) with BIC penalty ($\log(n)$) for each region. For $k = 1, \ldots, K_S$

(K_N), the covariate vector and its coefficients are modeled as

$$\begin{aligned}\boldsymbol{X}_{S,k}\boldsymbol{\beta} &= \beta_0 + \beta_1 \texttt{Extor}_k + \beta_2 \texttt{Shop}_k + \beta_3 \texttt{Street}_k + \beta_4 \texttt{Total}_k \\ &\quad + \beta_5 \texttt{Eco}_k + \beta_6 \texttt{Scholar}_k \qquad (6)\\ \boldsymbol{X}_{N,k}\boldsymbol{\beta} &= \beta_0 + \beta_2 \texttt{Shop}_k + \beta_4 \texttt{Total}_k + \beta_7 \texttt{Apart}_k \qquad (7)\end{aligned}$$

All covariates are centered and scaled. We use the same covariate vector for the LGCP model.

For sampling of $\boldsymbol{\beta}$ in the LGCP and NHPP, we implement an adaptive random walk MH algorithm (Andrieu and Thoms (2008)). We implement elliptical slice sampling for the GP in the LGCP (Murray and Adams (2010), Murray et al. (2010), Leininger (2014)). 20,000 samples are discarded as burn-in period and a subsequent 20,000 samples are preserved as posterior samples for the LGCP and the NHPP, respectively. Since, for spatial Gaussian processes, ϕ and σ^2 are not identifiable but the product, $\phi\sigma^2$ is (Zhang (2004)), we need to adopt an informative prior distribution for one of them. Here, we assume informative support for ϕ and adopt an inverse Gamma distribution for σ^2 with relatively large variance. As specific prior settings, we assume $\sigma^2 \sim \mathcal{IG}(2, 0.1)$ (inverse gamma), $\boldsymbol{\beta} \sim \mathcal{N}(\mathbf{0}, 100\mathbf{I})$ (normal) and $\phi \sim \mathcal{U}[0, 10]$ (uniform), where, after rescaling the eastings and northings, distances are in kilometers.

We also report the inefficiency factor for comparing sampling efficiency. The inefficiency factor (IF) is the ratio of the numerical variance of the estimate from the MCMC samples relative to that from hypothetical uncorrelated samples, and is defined as $1 + 2\sum_{s=1}^{\infty} \rho_s$ where ρ_s is the sample autocorrelation at lag s. It suggests the relative number of correlated draws necessary to attain the same variance for the posterior sample mean as from uncorrelated draws (e.g., Chib (2001)). Smaller values express more efficient sampling.

The estimation results are presented in Tables 1 and 2. The results for the covariates are similar for both the NHPP and LGCP though the posterior variances in the LGCP are larger than those in NHPP which is attributable to the additional randomness introduced by the Gaussian process. The reported log likelihood ("like" in the tables) for the LGCP is better (larger) than that for the NHPP in both regions.

We also implement model validation using (1) predictive interval coverage and (2) ranked probability score. Again, we set $p = 0.5$ for dividing into training and test datasets. Figure 4 shows the PIC with 90% nominal level and the RPS for both regions. Here, w denotes the number of randomly selected blocks for model comparison. As for the choice of B_r, since the total number of grid cells for this dataset is small, here we choose $w = 1, \ldots, 10$ grids from the 22 grids in the north and 68 grids in the south, rather than choosing B_r with respect to the rate q. The

Table 1: Estimation results of models for theft locations in the north region

	LGCP				NHPP			
	Mean	Stdev	95% Int	IF	Mean	Stdev	95% Int	IF
β_0	6.423	0.096	[6.226, 6.607]	194	6.434	0.042	[6.350, 6.514]	15
β_2	-0.250	0.104	[-0.438, -0.036]	273	-0.259	0.047	[-0.351, -0.167]	14
β_4	0.604	0.117	[0.341, 0.823]	250	0.556	0.049	[0.457, 0.651]	15
β_7	-0.119	0.084	[-0.275, 0.070]	191	-0.142	0.043	[-0.226, -0.057]	13
σ^2	0.081	0.049	[0.027, 0.218]	258				
ϕ	4.406	2.634	[0.363, 9.493]	730				
$\sigma^2\phi$	0.328	0.266	[0.031, 1.014]	391				
like	-68.47	3.185	[-75.46, -62.88]	64	-81.98	1.371	[-85.44, -80.28]	9

Table 2: Estimation results of models for theft locations in the south region

	LGCP				NHPP			
	Mean	Stdev	95% Int	IF	Mean	Stdev	95% Int	IF
β_0	8.071	0.048	[7.977, 8.178]	502	8.050	0.018	[8.014, 8.086]	38
β_1	-0.143	0.043	[-0.233, -0.064]	231	-0.117	0.024	[-0.164, -0.068]	35
β_2	-0.169	0.051	[-0.287, -0.077]	474	-0.157	0.025	[-0.206, -0.109]	9
β_3	-0.143	0.043	[-0.233, -0.060]	281	-0.147	0.025	[-0.196, -0.097]	20
β_4	0.624	0.072	[0.485, 0.768]	309	0.574	0.043	[0.489, 0.660]	13
β_5	-0.068	0.041	[-0.155, 0.013]	343	-0.076	0.023	[-0.121, -0.029]	17
β_6	0.169	0.037	[0.099, 0.243]	202	0.160	0.021	[0.118, 0.202]	22
σ^2	0.060	0.027	[0.032, 0.109]	217				
ϕ	3.202	2.300	[0.782, 9.114]	895				
$\sigma^2\phi$	0.180	0.131	[0.048, 0.540]	691				
like	-228.2	6.292	[-242.4, -216.7]	124	-294.3	1.794	[-298.6, -291.7]	22

LGCP outperforms the NHPP, more so for the south, the larger dataset.

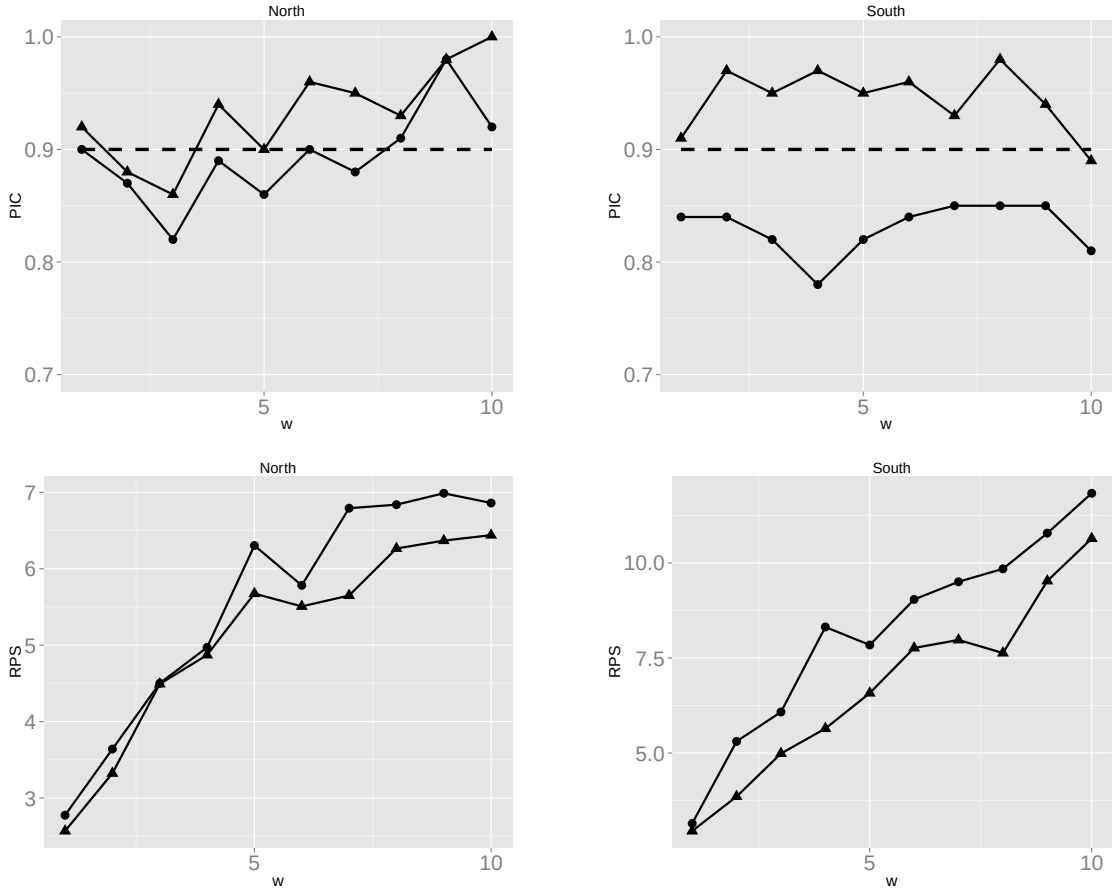

Figure 4: PIC (top) with 90% nominal level (dashed line) and RPS (bottom) for the north (left) and south (right) regions: NHPP (•) and LGCP (▲). w is the number of randomly selected blocks for model comparison.

4 A conditional model specification

4.1 Conditioning recovery location on theft location

For the analysis of recovery locations, we consider a conditional density specification given theft location. Now, we do not view the data as a point realization. We do not have to specify a set in which we consider our recovery locations; we can include some recovery points located outside the Neza region. Also, we do not have to split the Neza region for this analysis. We can allow the theft location to affect not only the mean for the recovery location but also the uncertainty in the recovery location.

We denote by $\boldsymbol{s}_R$ a recovery location and by $\boldsymbol{s}_T$ a theft location with $\mathcal{S}_T = \{\boldsymbol{s}_{T,1}, \ldots, \boldsymbol{s}_{T,n}\}$ and $\mathcal{S}_R = \{\boldsymbol{s}_{R,1}, \ldots, \boldsymbol{s}_{R,m}\}$ where $m < n$. We denote the conditional density specification for recovery location $\boldsymbol{s}_R$ given a theft location $\boldsymbol{s}_T$ as $f_R(\boldsymbol{s}_R|\boldsymbol{s}_T)$.

Combined with the marginal point pattern model for theft locations in the previous section, we create a joint model for theft location and recovery location. In this way, we can employ all of the theft data and all of the available recovery data. However, this model is not a *joint* specification in the sense of viewing the data as a point pattern of pairs of locations over a bounded set in $\mathbb{R}^2 \times \mathbb{R}^2$. We defer this model to the next section.

Let $\mathcal{S}_T^* = \{\boldsymbol{s}_{T,1}^*, \ldots, \boldsymbol{s}_{T,m}^*\}$ be the theft locations corresponding to recovery points, i.e., $\boldsymbol{s}_{T,j}^*$ is the corresponding theft location for the recovery point $\boldsymbol{s}_{R,j}$ for $j = 1, \ldots, m$. For $j = 1, \ldots, m$,

$$f_R(\boldsymbol{s}_{R,j}|\boldsymbol{s}_{T,j}^*) \propto |\Sigma(\boldsymbol{s}_{T,j}^*)|^{-1/2} \exp\left(-(\boldsymbol{s}_{R,j} - \boldsymbol{s}_{T,j}^*)'\Sigma(\boldsymbol{s}_{T,j}^*)^{-1}(\boldsymbol{s}_{R,j} - \boldsymbol{s}_{T,j}^*)\right), \quad (8)$$

$\Sigma(\boldsymbol{s}_{T,j}^*)$ is 2×2 covariance kernel dependent on theft location $\boldsymbol{s}_{T,j}^*$. A benchmark specification would assume a constant covariance kernel across theft locations, i.e.,

$$\Sigma(\boldsymbol{s}_{T,j}^*) = \Sigma = \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix} \quad (9)$$

A locally adaptive covariance kernel can be also considered, for example, employing the spatially varying covariance kernel in Higdon et al. (1999),

$$\begin{aligned} \Sigma(\boldsymbol{s}_T)^{\frac{1}{2}} =& \sigma \begin{pmatrix} \left(\frac{\sqrt{4A^2 + \|\psi(\boldsymbol{s}_T)\|^4\pi^2}}{2\pi} + \frac{\|\psi(\boldsymbol{s}_T)\|^2}{2}\right)^{\frac{1}{2}} & 0 \\ 0 & \left(\frac{\sqrt{4A^2 + \|\psi(\boldsymbol{s}_T)\|^4\pi^2}}{2\pi} - \frac{\|\psi(\boldsymbol{s}_T)\|^2}{2}\right)^{\frac{1}{2}} \end{pmatrix} \\ & \times \begin{pmatrix} \cos(\alpha(\boldsymbol{s}_T)) & \sin(\alpha(\boldsymbol{s}_T)) \\ -\sin(\alpha(\boldsymbol{s}_T)) & \cos(\alpha(\boldsymbol{s}_T)) \end{pmatrix} \end{aligned} \quad (10)$$

where $\|\psi(\boldsymbol{s}_T)\|^2 = \psi_x(\boldsymbol{s}_T)^2 + \psi_y(\boldsymbol{s}_T)^2$, $\alpha(\boldsymbol{s}_T) = \tan^{-1}\psi_y(\boldsymbol{s}_T)/\psi_x(\boldsymbol{s}_T)$ and $A = 3.5$ as fixed in Higdon et al. (1999). $\psi_x(\boldsymbol{s})$ and $\psi_y(\boldsymbol{s})$ are independent Gaussian processes with mean 0 and common Gaussian covariance function $C(\boldsymbol{s}_{T,i}, \boldsymbol{s}_{T,j}) = \exp(-\phi_*\|\boldsymbol{s}_{T,i} - \boldsymbol{s}_{T,j}\|^2)$. They introduce spatial dependence in $\Sigma(\boldsymbol{s}_T)$. ϕ_* is a tuning parameter which determines the spatial decay of the Gaussian processes. We fix this parameter at several different values in the ensuing analysis.

As a last remark here, there is no evident way to introduce spatial covariates such as those noted in the previous section into the conditional model. The mean for the recovery location should be the theft location; a regression specification here is not sensible. Furthermore, with the flexibility of a location dependent covariance matrix to accommodate direction and dispersion, we could not gain more flexibility by attempting to insert covariate information associated with $\boldsymbol{s}_T$ into $\Sigma(\boldsymbol{s}_T)$.

Model comparison using the bivariate CRPS

Let $\{\boldsymbol{s}_{T,h}^{test}, \boldsymbol{s}_{R,h}^{test}\}_{h=1}^{H}$ be the randomly selected test samples for evaluating predictive performance and $\{\boldsymbol{s}_{T,j}^{train}, \boldsymbol{s}_{R,j}^{train}\}_{j=1}^{m-H}$ be the remaining training samples for parameter estimation. For comparing predictive performance, we calculate the *bivariate* continuous rank probability score (CRPS) as discussed in Gneiting et al. (2008). In this regard, we value more a bivariate distribution $F(\cdot|\boldsymbol{s}_T^{test})$ if it is more concentrated around $\boldsymbol{s}_R^{test}$. The criterion is calculated through Monte Carlo integrations, using draws from $F(\cdot|\boldsymbol{s}_T^{test})$, as

$$CRPS(F(\cdot|\boldsymbol{s}_{T,h}^{test}), \boldsymbol{s}_{R,h}^{test}) = \frac{1}{L}\sum_{\ell=1}^{L} \|\boldsymbol{s}_{R,h}^{(\ell)} - \boldsymbol{s}_{R,h}^{test}\| - \frac{1}{2L^2}\sum_{\ell=1}^{L}\sum_{\ell'=1}^{L} \|\boldsymbol{s}_{R,h}^{(\ell)} - \boldsymbol{s}_{R,h}^{(\ell')}\| \qquad (11)$$

Here, $\boldsymbol{s}_{R,h}^{(\ell)}$ are samples from $F(\cdot|\boldsymbol{s}_T^{test})$. Fitting within a Bayesian framework enables posterior samples from $F(\cdot|\boldsymbol{s}_T^{test})$ using posterior samples of the model parameters. In particular, for construction of the bivariate predictive distribution with spatially varying kernel, given the ℓ-th posterior sample of $\psi_x(\boldsymbol{s})$, $\psi_y(\boldsymbol{s})$ at training samples $\{\boldsymbol{s}_{T,j}^{train}\}_{j=1}^{m-H}$ and σ, we generate $\psi_x(\boldsymbol{s})$, $\psi_y(\boldsymbol{s})$ at test samples $\{\boldsymbol{s}_{T,h}^{test}\}_{h=1}^{H}$ from the conditional Gaussian distribution and calculate $\Sigma(\boldsymbol{s}_{T,h}^{test})$ for $h = 1, \ldots, H$.

4.2 Results for the conditional density specification for recovery locations

For recovery locations, we implement conditional density specification with constant and spatially varying covariance kernels for both datasets. For the constant covariance kernel parameters, we assume $\sigma_1^2, \sigma_2^2 \sim \mathcal{IG}(2, 0.1)$ and $\rho \in \mathcal{U}[-1, 1]$. We discard the first 10,000 samples and retain the subsequent 10,000 samples as posterior samples. For the spatially varying covariance kernel parameter, we assume $\sigma^2 \sim \mathcal{IG}(2, 0.1)$. For this model, we consider three fixed ϕ_* values: (i) $\phi_* = 30$, (ii) $\phi_* = 10$ and (iii) $\phi_* = 1$. We discard the first 20,000 samples and retain the subsequent 20,000 samples as posterior samples.

Computation for the Neza dataset is manageable. However, the number of theft locations in Belo Horizonte is a bit large (in terms of matrix inversion and determinant calculation) to sample the Gaussian processes at all $\boldsymbol{s}_T$. So, we approximate the study region by using 305 disjoint regular grid cells. $\Sigma(\boldsymbol{s}_T)$ is evaluated at the nearest grid centroid. For sampling the Gaussian processes $\psi_x(\boldsymbol{s})$ and $\psi_y(\boldsymbol{s})$, we implement elliptical slice sampling. The estimation results are given in Table 3. The spatially varying covariance model fits better than the constant covariance kernel model with respect

to the loglikelihood, preferring the larger values of ϕ_* (weaker spatial dependence in the $\Sigma(\cdot)$'s) for the Neza, less so for Belo Horizonte.

Table 3: Estimation results of conditional model specification

	Neza				Belo Horizonte			
	Mean	Stdev	95% Int	IF	Mean	Stdev	95% Int	IF
Constant								
σ_1	2.360	0.084	[2.199, 2.542]	13	3.012	0.028	[2.957, 3.071]	13
σ_2	2.142	0.076	[2.001, 2.303]	9	3.953	0.038	[3.875, 4.028]	11
ρ	-0.421	0.041	[-0.502, -0.334]	8	0.039	0.013	[0.013, 0.067]	17
like	-961.9	1.214	[-965.0, -960.5]	14	-18269	1.231	[-18272, -18268]	12
Spatial								
(i) $\phi_* = 30$								
σ	1.527	0.050	[1.430, 1.628]	226	2.796	0.019	[2.757, 2.835]	45
like	-744.2	15.28	[-770.3, -711.4]	649	-16632	16.15	[-16666, -16603]	730
(ii) $\phi_* = 10$								
σ	1.532	0.049	[1.436, 1.631]	242	2.796	0.019	[2.756, 2.837]	32
like	-746.6	14.56	[-775.0, -717.0]	681	-16631	16.13	[-16660, -16598]	457
(iii) $\phi_* = 1$								
σ	1.693	0.047	[1.604, 1.788]	47	2.790	0.020	[2.752, 2.830]	36
like	-822.9	8.519	[-838.8, -805.7]	259	-16613	16.93	[-16656, -16586]	672

Lastly here, we compare model performance by calculating bivariate CRPS as defined in (11). For Neza (Belo Horizonte), we hold out randomly $H = 80$ (2625) samples for testing, and estimate parameter values for the remaining 302 (2625) training dataset. Then, given test theft locations $\{\boldsymbol{s}_{T,h}^{test}\}_{h=1}^{H}$, we calculate the predictive conditional densities for the corresponding recovery locations. That is, for the spatially varying kernel, we predict $\psi_x(\boldsymbol{s}_{T,h}^{test})$ and $\psi_y(\boldsymbol{s}_{T,h}^{test})$ given posterior samples σ and $\psi_x(\boldsymbol{s})$ and $\psi_y(\boldsymbol{s})$. For Neza, the estimated bivariate CRPS's for the 80 test pairs are 2.947 (Σ constant), 2.624 ($\phi_* = 30$), 2.636 ($\phi_* = 10$) and 2.962 ($\phi_* = 1$). The spatially varying kernel models with $\phi_* = 30$ and $\phi_* = 10$ show similar performance, being preferred to the spatially varying kernel model with $\phi_* = 1$ and constant kernel model. For Belo Horizonte, we hold out randomly $H = 2,625$ (50% of the total number of points). The estimated bivariate CRPS for 2,625 test pairs are 156.43 (Σ constant), 152.59 ($\phi_* = 30$), 152.29 ($\phi_* = 10$) and 152.64 ($\phi_* = 1$). The spatially varying kernel models are indistinguishable and slightly outperform the constant kernel model.

5 Joint Point Pattern Modeling

The LGCP for pairs of locations

Now, we introduce joint point process models over $D \times D \subset \mathbb{R}^2 \times \mathbb{R}^2$. We denote observed pairs as $\mathcal{S}_P = \{\boldsymbol{s}_{P,1}, \ldots, \boldsymbol{s}_{P,m}\} = \{(\boldsymbol{s}_{R,1}, \boldsymbol{s}_{T,1}^*), \ldots, (\boldsymbol{s}_{R,m}, \boldsymbol{s}_{T,m}^*)\}$. The

intensity function for observed pairs is defined as

$$\log \lambda(\boldsymbol{s}_R, \boldsymbol{s}_T^*) = \boldsymbol{X}_R(\boldsymbol{s}_R)\boldsymbol{\beta}_R + \boldsymbol{X}_T(\boldsymbol{s}_T^*)\boldsymbol{\beta}_T + \eta(\boldsymbol{s}_R - \boldsymbol{s}_T^*)'\Sigma(\boldsymbol{s}_T^*)^{-1}(\boldsymbol{s}_R - \boldsymbol{s}_T^*) + z_R(\boldsymbol{s}_R) + z_T(\boldsymbol{s}_T^*), \quad (12)$$

$$\boldsymbol{z}_R \sim \mathcal{N}(\boldsymbol{0}, \mathbf{C}_{\boldsymbol{z}_R}), \quad \boldsymbol{z}_T \sim \mathcal{N}(\boldsymbol{0}, \mathbf{C}_{\boldsymbol{z}_T}). \quad (13)$$

Here, $z_R(\boldsymbol{s})$ and $z_T(\boldsymbol{s})$ are mean 0 GP's with covariance functions C_R and C_T, respectively. $\boldsymbol{z}_R = (z(\boldsymbol{s}_{R,1}), \ldots, z(\boldsymbol{s}_{R,m}))$ $(\boldsymbol{z}_T = (z(\boldsymbol{s}_{T,1}^*), \ldots, z(\boldsymbol{s}_{T,m}^*)))$ and $\mathbf{C}_{\boldsymbol{z}_R} = (C_R(\boldsymbol{s}_{R,i}, \boldsymbol{s}_{R,j}))_{i,j=1,\ldots,m}$ $(\mathbf{C}_{\boldsymbol{z}_T} = (C_T(\boldsymbol{s}_{T,i}^*, \boldsymbol{s}_{T,j}^*))_{i,j=1,\ldots,m})$. We assume the exponential covariance functions for C_R and C_T, i.e., $C_R(\boldsymbol{u}, \boldsymbol{u}') = \sigma_R^2 \exp(-\phi_R \|\boldsymbol{u} - \boldsymbol{u}'\|)$ and $C_T(\boldsymbol{u}, \boldsymbol{u}') = \sigma_T^2 \exp(-\phi_T \|\boldsymbol{u} - \boldsymbol{u}'\|)$. The first two terms of the log intensity introduce origin and destination covariates, the third term provides a local (spatially varying) distance measure between the origin location and the destination location, $\boldsymbol{s}_R$ and $\boldsymbol{s}_T^*$ through $\Sigma(\boldsymbol{s}_T^*)$, the spatially varying kernel presented in the previous section. η is expected to be negative, i.e., the recovery locations are more observed near the corresponding theft locations. In addition, the local $\Sigma(\boldsymbol{s}_T^*)$ enables directional preference for $\boldsymbol{s}_R$ near the boundaries of D. The fourth and fifth terms provide origin and destination spatial random effects using Gaussian processes. Without them, we have the analogue of an NHPP; with them, we have the analogue of a LGCP. However, we will need a large number of origin-destination pairs in order to learn about these local random effects adjustments. This joint specification is evidently richer than the marginal model with the conditional version from the previous section. However, it only employs the complete pairs in the data.

As above, the likelihood is approximated by gridding D into K blocks. Now, employing $K \times K$ blocks for $D \times D$, we obtain

$$\mathcal{L}(\mathcal{S}_P) \propto \exp\left(-\int_D \int_D \lambda(\boldsymbol{u}_R, \boldsymbol{u}_T) d\boldsymbol{u}_T d\boldsymbol{u}_R\right) \prod_{j=1}^{m} \lambda(\boldsymbol{s}_{R,j}, \boldsymbol{s}_{T,j}^*) \quad (14)$$

$$\approx \exp\left(-\sum_{k=1}^{K}\sum_{k'=1}^{K} \lambda(\boldsymbol{u}_{R,k}, \boldsymbol{u}_{T,k'}) \Delta_{T,k} \Delta_{T,k'}\right) \prod_{k=1}^{K}\prod_{k'=1}^{K} \lambda(\boldsymbol{u}_{R,k}, \boldsymbol{u}_{T,k'})^{n_{kk'}} \quad (15)$$

where $\sum_{k=1}^{K}\sum_{k'=1}^{K} n_{kk'} = m$. Note that $\boldsymbol{u}_{R,k} = \boldsymbol{u}_{T,k}$ for $k = 1, \ldots, K$.

p-thinning cross validation for the joint intensity model

For this joint intensity model, we sample evaluation grid cells $B_{R,r} \times B_{T,r} \subset D \times D$.

Then, we calculate the RPS, via Monte Carlo integration, using posterior samples, as

$$RPS(F, N^{test}(B_{R,r} \times B_{T,r})) = \frac{1}{L}\sum_{\ell=1}^{L} |N^{(\ell)}(B_{R,r} \times B_{T,r}) - N^{test}(B_{R,r} \times B_{T,r})| - \frac{1}{2L^2}\sum_{\ell=1}^{L}\sum_{\ell'=1}^{L} |N^{(\ell)}(B_{R,r} \times B_{T,r}) - N^{(\ell')}(B_{R,r} \times B_{T,r})| \tag{16}$$

In our analysis, we choose $B_{R,r}$ and $B_{T,r}$ from K blocks randomly, so we have $K \times K$ possible combinations of $(B_{R,r}, B_{T,r})$.

5.1 Results

Results for the smaller dataset with covariates in Neza

Since 56 recovery points are outside the Neza region, we delete these points in this analysis. Furthermore, the number of observed pairs in the north region is only 32, and there are only 7 *cross pairs*, i.e., theft in the north and recovery in the south or theft in the south and recovery in the north. So, in this analysis we focus only on the 287 observed pairs in the south region (having ($K = 68$) grid cells). We fit four models: (i) an NHPP without spatial covariates (NHPP-Ind), (ii) an NHPP with spatial covariates (NHPP-Dep), (iii) a LGCP without spatial covariates (NHPP-Ind) and (iv) a LGCP with spatial covariates (LGCP-Dep).

For covariate selection, we start with NHPP-Ind model. As in Section 3, we implement the forward and backward selection separately for the recovery locations and for the theft locations. Two covariates are selected yielding

$$\boldsymbol{X}_{R,k}\boldsymbol{\beta} = \beta_0 + \beta_{R,2}\texttt{Shop}_k + \beta_{R,4}\texttt{Total}_k \tag{17}$$

$$\boldsymbol{X}_{T,k}\boldsymbol{\beta} = \beta_{T,2}\texttt{Shop}_k + \beta_{T,4}\texttt{Total}_k \tag{18}$$

We employ these same covariates for the other models. For identifiability, the LGCP below is just an LGCP for recovery locations, i.e., including only $z_R(\boldsymbol{s})$.

For sampling of $\boldsymbol{\beta}$ in LGCP and NHPP, we implement an adaptive random walk MH algorithm. We implement elliptical slice sampling for $\psi_x(\boldsymbol{s})$, $\psi_y(\boldsymbol{s})$ and $z_R(\boldsymbol{s})$. We fix the tuning parameter $\phi_* = 1$. We discard 20,000 samples as burn-in period and retain the subsequent 20,000 samples as posterior samples for the LGCP and the NHPP. As with the LGCP for theft locations, we assume $\sigma_R^2 \sim \mathcal{IG}(2, 0.1)$, $\boldsymbol{\beta} \sim \mathcal{N}(\mathbf{0}, 100\mathbf{I})$ and $\phi_R \sim \mathcal{U}[0, 10]$. The estimation results are shown in Tables 4 and 5. The estimation results for the covariates are similar in the NHPP and the LGCP though the posterior variances in the LGCP are larger than those in the NHPP. The

negative value of η suggests that the recovery locations tend to be located to near the theft locations. The estimated log likelihood is larger for the dependent models than for the independent models.

Table 4: Estimation results of NHPP for pairs: Neza

	NHPP-Ind				NHPP-Dep			
	Mean	Stdev	95% Int	IF	Mean	Stdev	95% Int	IF
β_0	5.535	0.067	[5.399, 5.665]	17	6.690	0.099	[6.493, 6.882]	71
$\beta_{R,2}$	-0.196	0.087	[-0.374, -0.028]	7	-0.171	0.085	[-0.345, -0.006]	13
$\beta_{R,4}$	0.238	0.083	[0.072, 0.401]	8	0.071	0.085	[-0.093, 0.241]	16
$\beta_{T,2}$	-0.118	0.075	[-0.270, 0.027]	9	-0.106	0.077	[-0.258, 0.045]	18
$\beta_{T,4}$	0.441	0.085	[0.277, 0.615]	7	0.319	0.088	[0.148, 0.495]	23
η					-0.128	0.016	[-0.160, -0.097]	95
like	-1067	1.621	[-1071, -1065]	11	-933	8.548	[-952, -917]	339

Table 5: Estimation results of LGCP for pairs: Neza

	LGCP-Ind				LGCP-Dep			
	Mean	Stdev	95% Int	IF	Mean	Stdev	95% Int	IF
β_0	5.600	0.203	[5.234, 6.091]	260	6.989	0.279	[6.482, 7.619]	449
$\beta_{R,2}$	-0.244	0.105	[-0.451, -0.040]	99	-0.228	0.111	[-0.448, -0.010]	94
$\beta_{R,4}$	0.313	0.118	[0.078, 0.555]	79	0.209	0.138	[-0.056, 0.488]	147
$\beta_{T,2}$	-0.117	0.075	[-0.266, 0.031]	9	-0.109	0.079	[-0.267, 0.046]	18
$\beta_{T,4}$	0.400	0.085	[0.277, 0.605]	5	0.343	0.090	[0.164, 0.523]	14
σ_R^2	0.221	0.108	[0.082, 0.498]	307	0.316	0.157	[0.109, 0.744]	312
ϕ_R	1.086	0.810	[0.211, 3.480]	842	0.928	0.558	[0.205, 2.323]	575
$\phi_R\sigma_R^2$	0.227	0.203	[0.042, 0.806]	731	0.257	0.148	[0.066, 0.634]	310
η					-0.143	0.017	[-0.176, -0.110]	129
like	-1043	4.980	[-1053, -1033]	185	-900	9.527	[-919, -882]	240

We implement model validation using (i) predictive interval coverage and (ii) ranked probability scores. Since the number of observed pairs is relatively small, we compare in-sample performance. The left side of Figure 5 shows PIC with 90% nominal level while the right side shows the RPS score. We select $q = 0.01, \ldots, 0.1$ where $q|D| \times |D|$. Although the differences among the models are small (likely due to the small size of the point pattern), the LGCP-Dep model is the best performer.

Results for the larger dataset without covariates in Belo Horizonte

The number of observed theft-recovery pairs for Belo Horizonte is 5,250. Without covariates, the intensity for Belo Horizonte is

$$\log \lambda(\boldsymbol{s}_R, \boldsymbol{s}_T^*) = \beta_0 + \eta(\boldsymbol{s}_R - \boldsymbol{s}_T^*)'\Sigma(\boldsymbol{s}_T^*)^{-1}(\boldsymbol{s}_R - \boldsymbol{s}_T^*) + z_R(\boldsymbol{s}_R) + z_T(\boldsymbol{s}_T^*), \tag{19}$$

$$\boldsymbol{z}_R(\mathcal{S}_R) \sim \mathcal{N}(\boldsymbol{0}, \mathbf{C}_R(\mathcal{S}_R, \mathcal{S}_R)), \quad \boldsymbol{z}_T(\mathcal{S}_T^*) \sim \mathcal{N}(\boldsymbol{0}, \mathbf{C}_T(\mathcal{S}_T^*, \mathcal{S}_T^*)). \tag{20}$$

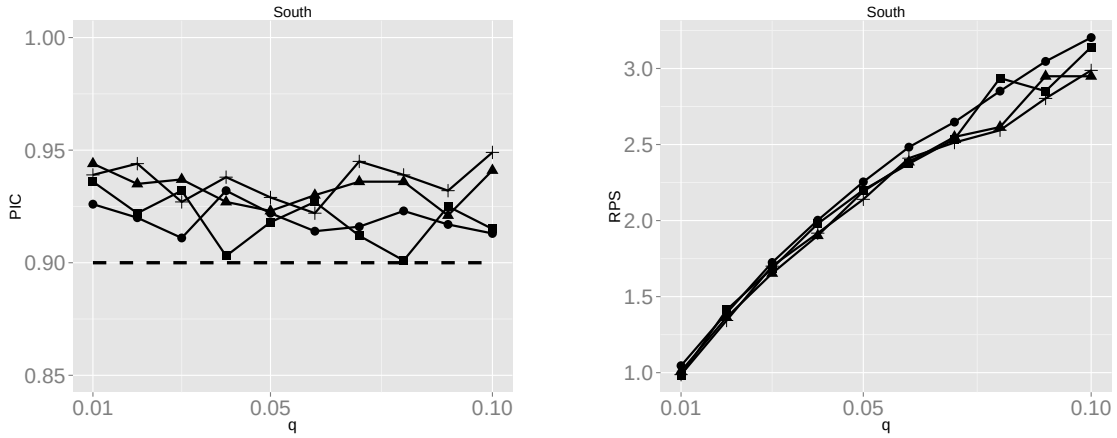

Figure 5: In-sample PIC (left) with 90% nomial level (dashed line) and RPS (right): NHPP-Ind (•), NHPP-Dep (▲), LGCP-Ind (■) and LGCP-Dep (+). q is the relative size of the evaluation grid B, i.e., $|B| = q|D|^2$.

We take $K = 305$ grids, the as same as in conditional density specification. Without covariates, we estimate two models: (i) a LGCP without the spatially varying distance measure (LGCP-Ind, i.e., $\eta = 0$) and (ii) a LGCP with this measure (LGCP-Dep). For sampling β_0, we implemented an adaptive random walk MH algorithm. We implemented elliptical slice sampling for $\psi_x(\boldsymbol{s})$, $\psi_y(\boldsymbol{s})$, $z_R(\boldsymbol{s})$ and $z_T(\boldsymbol{s})$. We fixed the tuning parameter $\phi_* = 1$. We discard 20,000 samples as the burn-in period and retain the subsequent 20,000 samples as posterior samples. We assume $\sigma_R^2, \sigma_T^2 \sim \mathcal{IG}(2, 0.1)$, $\beta_0 \sim \mathcal{N}(0, 100)$ and $\phi_R, \phi_T \sim \mathcal{U}[0, 10]$. Table 6 shows the estimation results for the LGCP models. The likelihood value for the LGCP-Dep is much larger than that of the LGCP-ind. The estimated value of η is negative (with small variance) as we expect.

Since the Belo Horizonte dataset has larger number of pairs, we investigate out of sample predictive performance by comparing PIC and RPS. We use 2,625 pairs for the estimation of parameters and intensity surface and remaining 2,625 as test pairs. The left side of Figure 6 shows PIC with 90% nominal level and the right side shows RPS scores. We select $q = 0.02, \ldots, 0.2$ where $q|D| \times |D|$. The LGCP-dep model shows a bit of under-coverage on the PIC but outperforms the LGCP-ind model both on PIC and RPS.

6 A ridge model specification

6.1 A connection with trend surface modeling

We conclude our analyses with a brief digression which looks at the point pattern of say theft locations in a different way. In particular, for a point pattern, $\mathcal{S}$ over

Table 6: Estimation results of LGCP for pairs: Belo Horizonte

	LGCP-Ind				LGCP-Dep			
	Mean	Stdev	95% Int	IF	Mean	Stdev	95% Int	IF
β_0	7.780	0.232	[7.269, 8.208]	1223	10.41	0.249	[10.07, 10.80]	1320
σ^2_R	3.170	0.721	[2.289, 4.730]	1406	2.354	0.339	[1.840, 3.176]	979
ϕ_R	0.203	0.054	[0.125, 0.290]	1421	0.257	0.030	[0.192, 0.315]	830
$\phi_R\sigma^2_R$	0.608	0.063	[0.519, 0.781]	1020	0.604	0.099	[0.472, 0.826]	1244
σ^2_T	3.049	0.665	[2.018, 4.442]	1361	1.553	0.396	[1.072, 2.340]	1205
ϕ_T	0.148	0.034	[0.097, 0.227]	1319	0.289	0.082	[0.151, 0.413]	1343
$\phi_T\sigma^2_T$	0.433	0.040	[0.359, 0.513]	949	0.420	0.061	[0.314, 0.529]	1242
η					-0.044	0.001	[-0.046, -0.043]	206
like	-14348	15.61	[-14381, -14319]	934	-10789	21.06	[-10829, -10744]	750

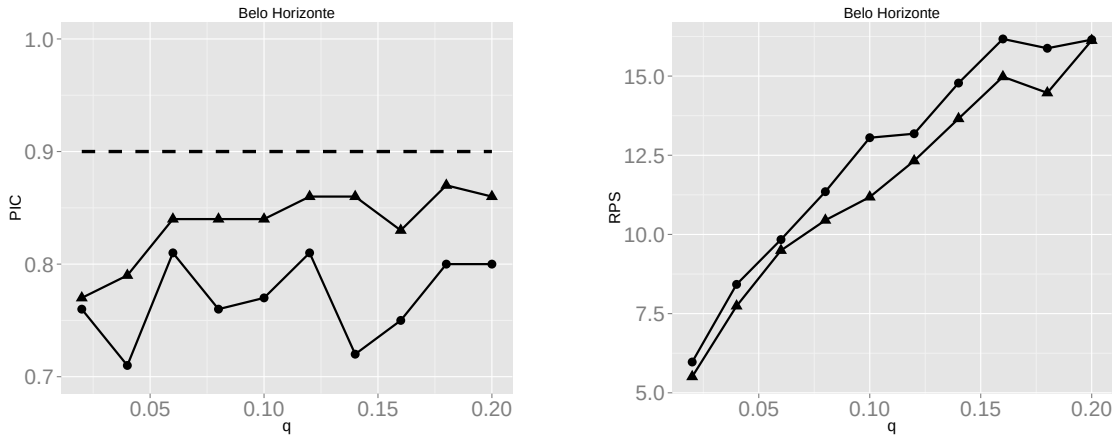

Figure 6: Out of sample PIC (left) with 90% nomial level (dashed line) and RPS (right): LGCP-Ind (•) and LGCP-Dep (▲). q is the relative size of the evaluation grid B, i.e., $|B| = q|D|^2$.

a set $D \subset \mathbb{R}^2$, let $(x(\boldsymbol{s}), y(\boldsymbol{s}))$ denote the x, y coordinates for point $\boldsymbol{s} \in \mathcal{S}$. The analogue of our joint origin-destination modeling in the previous section could be viewed as specifying an intensity over D of the form $\lambda(\boldsymbol{s}) = \lambda_0 e^{-\beta(y(\boldsymbol{s})-x(\boldsymbol{s}))^2}$ where β is expected to be larger than 0. Such a specification expects points close to the 45^o line. Illustratively, we could imagine D being the square $[-a, a] \times [-a, a]$, although we need not insist on a product set. We could add covariates to the intensity, $W_1(x(\boldsymbol{s}))$, $W_2(y(\boldsymbol{s}))$, even $W(\boldsymbol{s})$.

Note that this specification is a special case of a quadratic trend surface for $\log\lambda(\boldsymbol{s})$ over $\boldsymbol{s}$, i.e., $\beta_0 + \beta_1 x(\boldsymbol{s}) + \beta_2 y(\boldsymbol{s}) + \beta_3 x^2(\boldsymbol{s}) + \beta_4 y^2(\boldsymbol{s}) + \beta_5 x(\boldsymbol{s}) y(\boldsymbol{s})$. The sub-model investigates the adequacy of an intensity having a "ridge" along the 45^o line; points such that $|x(\boldsymbol{s}) - y(\boldsymbol{s})|$ is small are encouraged. In this regard, consider a more general $\lambda(\boldsymbol{s}) = \lambda_0 e^{-\beta(y(\boldsymbol{s})-(a+bx(\boldsymbol{s})))^2}$. Such a linear ridge in the intensity which might be appropriate when say, events with locations close to a relatively straight road are encouraged. Again, this model is a sub-model of the more general quadratic trend surface. Model comparison between this reduced model and the full quadratic would be useful if evidence of such *linear* behavior is seen in the point pattern.

Extending the quadratic surface, we could enrich β to $\beta(\boldsymbol{s})$ a GP over $\mathbb{R}^2$, allowing local weighting of the distances between $y(\boldsymbol{s})$ and $x(\boldsymbol{s})$. Alternatively, we could attempt local slopes and intercepts, in the spirit of Gelfand et al. (2003). That is, $\lambda(\boldsymbol{s}) = \lambda_0 e^{-\beta(y(\boldsymbol{s})-(a(\boldsymbol{s})+b(\boldsymbol{s})x(\boldsymbol{s})))^2}$ where $a(\boldsymbol{s})$ and $b(\boldsymbol{s})$ are a bivariate GP, a locally linear specification. This could handle a very flexible relationship between the $x(\boldsymbol{s})$ and the $y(\boldsymbol{s})$ but may be difficult to fit because the model will likely be too flexible for the data. Finally, connecting to the previous section, there, $(\boldsymbol{s}_O - \boldsymbol{s}_D)^T B(\boldsymbol{s}_O - \boldsymbol{s}_D)$ is the analogue of the distance between the x and y coordinates above. Furthermore, there B plays the role of β here, and enriching there B to $B(\boldsymbol{s}_O)$ is analogous to enriching the β's above.

6.2 A simulation example for points in the unit square

We develop a simulation considering two types of models: (1) linear ridge (LR) and (2) locally varying dispersion (LVD). That is,

$$\text{LR:} \quad \lambda(\boldsymbol{s}) = \lambda_0 \exp(-\beta(y(\boldsymbol{s}) - (a + bx(\boldsymbol{s})))^2) \tag{21}$$

$$\text{LVD:} \quad \lambda(\boldsymbol{s}) = \lambda_0 \exp(-\beta(\boldsymbol{s})(y(\boldsymbol{s}) - (a + b(\boldsymbol{s})x(\boldsymbol{s})))^2) \tag{22}$$

where $b(\boldsymbol{s})$ is a Gaussian process. First, we fit two cases of LR models: (case 1) $(\lambda_0, \beta, a, b) = (2000, 10, 0, 1)$ and (case 2) $(\lambda_0, \beta, a, b) = (4000, 100, 0, 1)$. The left and center panels of Figure 7 show the plots of the simulated points of LR models in the case 1 and 2. Case 2 shows stronger concentration around the 45^o line. We adopt

weak priors: $\lambda_0 \sim \mathcal{G}(2, 0.001)$, $\beta \sim \mathcal{N}(0, 500)$, $a \sim \mathcal{N}(0, 100)$ and $b \sim \mathcal{N}(0, 100)$. We discard 20,000 samples as burn-in period and retain the subsequent 20,000 samples as posterior samples.

The quadratic trend surface (TS) model for $\log \lambda(\boldsymbol{s})$ over $\boldsymbol{s}$ arising from expansion of LR model is

$$\begin{aligned}\beta(y(\boldsymbol{s}) - (a + bx(\boldsymbol{s})))^2 = \beta a^2 + 2ab\beta x(\boldsymbol{s}) - 2a\beta y(\boldsymbol{s}) \\ + \beta b^2 x^2(\boldsymbol{s}) + \beta y^2(\boldsymbol{s}) - 2b\beta x(\boldsymbol{s}) y(\boldsymbol{s}).\end{aligned} \tag{23}$$

We have the coefficients $(\beta_0, \beta_1, \beta_2, \beta_3, \beta_4, \beta_5)$ as functions of (λ_0, β, a, b), i.e.,

$$\beta_0 = \log(\lambda_0) - \beta a^2, \beta_1 = -2ab\beta, \beta_2 = 2a\beta, \beta_3 = -\beta b^2, \beta_4 = -\beta, \beta_5 = 2b\beta \tag{24}$$

We fit the simulated point patterns from LR models under cases 1 and 2 using LR and the full quadratic trend surface specification to investigate the recovery of parameter values. Table 7 shows the estimation results for the simulated point patterns from LR models for both fitting models under both cases. The LR and TS models recover the true parameter values. However, with the LR model being true, the LR fit shows less uncertainty, e.g., in estimating λ_0.

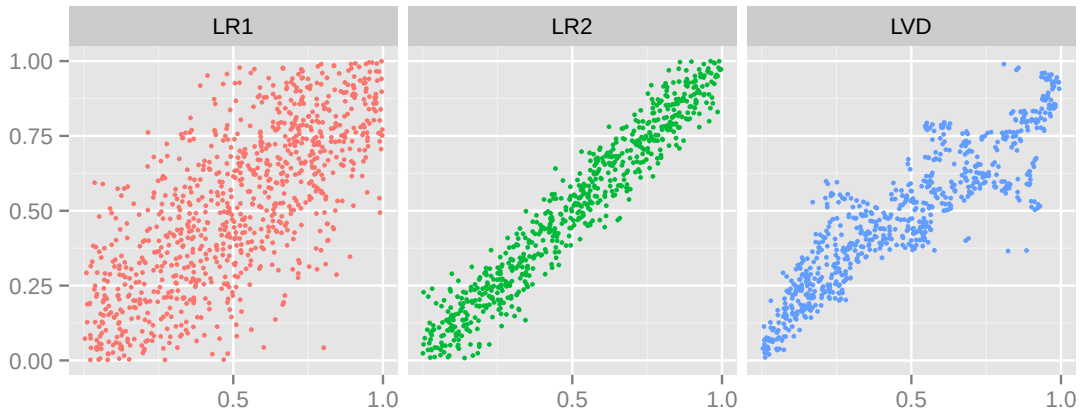

Figure 7: Simulated points from linear ridge models in case 1 (left) and case 2 (center) from locally varying dispersion model (right)

We also consider the locally varying dispersion model. We generate Gaussian process surfaces for $\log \beta(\boldsymbol{s}) \sim \mathcal{N}(\mu_\beta, C_\beta)$ and $b(\boldsymbol{s}) \sim \mathcal{N}(\mu_b, C_b)$. We assume the exponential covariance functions for C_β and C_b, i.e., $C_\beta(\boldsymbol{u}, \boldsymbol{u}') = \sigma_\beta^2 \exp(-\phi_\beta \|\boldsymbol{u} - \boldsymbol{u}'\|)$ and $C_b(\boldsymbol{u}, \boldsymbol{u}') = \sigma_b^2 \exp(-\phi_b \|\boldsymbol{u} - \boldsymbol{u}'\|)$. The true values are $(\lambda_0, \mu_\beta, a, \mu_b, \sigma_b^2, \phi_b, \sigma_\beta^2, \phi_\beta) = (6000, 4, 0, 1, 1, 1, 1, 1)$. The right panel of Figure 7 shows the realization of the simulated points. Evidence of a ridge still emerges but the locally varying slope di-

Table 7: Estimation results for simulated points from linear ridge

	Case 1				Case 2			
	True	Mean	95% Int	IF	True	Mean	95% Int	IF
LR								
λ_0	2000	1977	[1826, 2132]	32	4000	4037	[3684, 4409]	15
β	10	9.970	[8.704, 11.38]	148	100	99.31	[88.79, 110.6]	39
a	0	0.005	[-0.033, 0.051]	109	0	0.005	[-0.006, 0.017]	69
b	1	0.993	[0.907, 1.065]	120	1	0.990	[0.970, 1.011]	72
TS								
$\exp(\beta_0)$	2000	1772	[1356, 2295]	189	4000	3762	[2927, 4782]	72
β_1	0	0.000	[-1.076, 1.082]	63	0	-1.036	[-3.453, 1.371]	35
β_2	0	0.360	[-0.694, 1.408]	85	0	1.327	[-1.033, 3.808]	50
β_3	-10	-10.11	[-11.65, -8.640]	36	-100	-98.35	[-109.6, -87.36]	29
β_4	-10	-10.39	[-11.87, -8.870]	44	-100	-100.2	[-111.6, -88.92]	19
β_5	20	20.39	[18.16, 22.67]	30	200	198.4	[176.9, 220.7]	24

minishes the emphasis on the 45^o line. We assume weak priors: $\lambda_0 \sim \mathcal{G}(2, 0.001)$, $a, \mu_\beta, \mu_b \sim \mathcal{N}(0, 100)$, $\sigma^2_\beta, \sigma^2_b \sim \mathcal{IG}(2, 0.1)$, $\phi_\beta, \phi_b \sim \mathcal{U}[0, 10]$. We discard 30,000 samples as burn-in period and retain the subsequent 30,000 samples as posterior samples. Table 8 presents the estimation result of the locally varying dispersion model. We recover the parameter values, albeit with relatively larger variances.

Table 8: Estimation results for locally varying dispersion

	True	Mean	Stdev	95% Int	IF
LVD					
λ_0	6000	5452	309.3	[4885, 6104]	316
μ_β	4	4.532	0.366	[3.880, 5.290]	1669
a	0	0.014	0.018	[-0.019, 0.047]	1393
μ_b	1	0.965	0.039	[0.889, 1.033]	1344
σ^2_b	1	3.754	1.262	[1.542, 6.577]	1510
ϕ_b	1	0.263	0.092	[0.173, 0.484]	1641
$\sigma^2_b\phi_b$	1	0.923	0.290	[0.571, 1.744]	1630
σ^2_β	1	2.744	0.761	[1.620, 4.518]	1074
ϕ_β	1	7.256	1.806	[3.588, 9.911]	1813
$\sigma^2_\beta\phi_\beta$	1	0.741	0.407	[0.343, 1.945]	1363

7 Summary and Future Work

We have considered a little studied problem for point patterns namely the setting where we have a point pattern of origins over $D \subset \mathbb{R}^2$ (in our case, locations of car thefts) and an associated point pattern of destinations, again over $D \subset \mathbb{R}^2$ (in our case locations of car recoveries). We have presented a marginal approach for modeling the origins using NHPP's and LGCPs along with a conditional approach for modeling

destinations given origins. We have also considered a joint modeling approach where we view the point pattern as an origin-destination pair and specify a model over a subset of $\mathbb{R}^2 \times \mathbb{R}^2$. Making a connection to this last model, we have briefly considered looking at the set of locations as arising from a special case of a quadratic trend surface which encourages a ridge in the intensity.

Future work will investigate a much different application. We will investigate economic labor force data where, for an individual, we have the location where she/he resides as well as the location where she/he works. Working with metropolitan areas will provide much larger point patterns with much more demanding model fitting. Future work with theft-recovery data would introduce consideration of time, i.e., we will have not only the location of the theft but also the time of the theft. Similarly, we have not only a location for the recovery but as well the time of the recovery, with an implicit order in time for the latter relative to the former. Unfortunately, at present, neither of the datasets provide sufficient time information to enable such investigation.

Acknowledgements

This work partially funded by Grant MTM2016-78917-R from the Spanish Ministry of Science and Education, and Grant P1-1B2015-40 from University Jaume I. The work of Shinichiro Shirota was supported in part by the Nakajima Foundation. The authors thank Renato Assunção for providing the datasets and shape file of Belo Horizonte. The authors also thank the Nezahualcoyotl Town Hall, Arturo Arango, and Direccion General de Seguridad Ciudadana (in Mexico) for providing the Neza dataset. The computational results are obtained by using Ox version 7.1 (Doornik (2007)).

References

Adrienko, N. and G. Adrienko (2011). Spatial generalization and aggregation of massive movement data. *IEEE Transactions on Visualization and Computer Graphics 17*, 205–219.

Andrieu, C. and J. Thoms (2008). A tutorial on adaptive MCMC. *Statistics and Computing 18*, 343–373.

Assunção, R. M. and D. L. Lopes (2007). Testing association between origin-destination spatial locations. In C. A. Davis Jr and A. M. V. Monteiro (Eds.), *Advances in Geoinformatics*, pp. 293–304. Berlin: Springer.

Baddeley, A. and R. Turner (2000). Practical maximum pseudolikelihood for spatial point patterns. *Aust N Z J Stat 42*, 283–315.

Baddeley, A. and R. Turner (2005). `spatstat`: an R package for analyzing spatial point patterns. *Journal of Statistical Software 12*, 1–42.

Baddeley, A., R. Turner, J. Mateu, and A. Bevan (2013). Hybrids of Gibbs point process models and their implementation. *Journal of Statistical Software 55*, 1–43.

Banerjee, S., B. P. Carlin, and A. E. Gelfand (2014). *Hierarchical Modeling and Analysis for Spatial Data, 2nd ed.* Chapman and Hall/CRC.

Banerjee, S., A. E. Gelfand, and W. Polasek (2000). Geostatistical modelling of spatial interaction data with application to postal service performance. *Journal of Statistical Planning and Inference 90*, 87–105.

Beneš, V., K. Bodlák, J. Møller, and R. Waagepetersen (2005). A case study on point process modelling in disease mapping. *Image Analysis and Stereology 24*, 159–168.

Chakraborty, A., M. A. Beamonte, A. E. Gelfand, M. P. Alonso, P. Gargallo, and M. Salvador (2013). Spatial interaction models with individual-level data for explaining labor flows and developing local labor markets. *Computational Statistics and Data Analysis 58*, 292–307.

Chib, S. (2001). Markov chain Monte Carlo methods: computation and inference. In G. Elliott, C. W. J. Granger, and A. Timmermann (Eds.), *Handbook of Econometrics*, Volume 5, pp. 3569–3649. Amsterdam: North Holland Press.

Czado, C., T. Gneiting, and L. Held (2009). Predictive model assessment for count data. *Biometrics 65*, 1254–1261.

Gelfand, A. E., H. J. Kim, C. F. Sirmans, and S. Banerjee (2003). Spatial modeling with spatially varying coefficient processes. *Journal of the American Statistical Association 98*, 387–396.

Gneiting, T. and A. E. Raftery (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association 102*, 359–378.

Gneiting, T., L. I. Stanberry, E. P. Grimit, L. Held, and N. A. Johnson (2008). Assessing probabilitic forecasts of multivariate quantiles with applications to ensemble predictions of surface winds. *Test 17*, 211–235.

Guo, D., X. Zhu, H. Jin, P. Gao, and C. Andris (2012). Discovering spatial patterns in origin-destination mobility data. *Transactions in GIS 16*, 411–429.

Higdon, D., J. Swall, and J. Kern (1999). Non-stationary spatial modeling. *Bayesian Statistics 6*, 761–768.

Huang, F. and Y. Ogata (1999). Improvements of the maximum pseudo-likelihood estimators in various spatial statistical models. *Journal of Computational and Graphical Statistics 8*, 510–530.

Illian, J., A. Penttinen, H. Stoyan, and D. Stoyan (2008). *Statistical Analysis and Modelling of Spatial Point Patterns.* Wiley.

Leininger, T. J. (2014). Bayesian analysis of spatial point patterns. Dissertation.

Leininger, T. J. and A. E. Gelfand (2016). Bayesian inference and model assessment for spatial point patterns using posterior predictive samples. *Bayesian Analysis*. forthcomming.

Lopes, D. and R. Assunção (2012). Visualizing marked spatial and origin-destination point patterns with dynamically linked windows. *Journal of Computational and Graphical Statistics 21*, 134–154.

Møller, J., A. R. Syversveen, and R. P. Waagepetersen (1998). Log Gaussian Cox processes. *Scandinavian Journal of Statistics 25*, 451–482.

Møller, J. and R. Waagepetersen (2004). *Statistical Inference and Simulation for Spatial Point Processes.* Chapman and Hall/RC.

Murray, I. and R. P. Adams (2010). Slice sampling covariance hyperparameters of latent Gaussian models. In *Advances in Neural Information Processing Systems 23*, Cambridge, MA. MIT Press.

Murray, I., R. P. Adams, and M. M. Graham (2010). Elliptical slice sampling. In *Proceedings of the 13th International Conference on Artifical Intelligence and Statistics (AISTAT).* AISTAT Press.

Robert, C. P. and G. Casella (2004). *Monte Carlo Statistical Methods.* Springer-Verlag.

Waagepetersen, R. (2004). Convergence of posteriors for discretized log Gaussin Cox processes. *Stochastic and Probability Letters 66*, 229–235.

Wood, J., J. Dykes, and A. Slingsby (2010). Visualisation of origins, destinations and flows with OD maps. *Cartographic Journal 47*, 117–129.

Zhang, H. (2004). Inconsistent estimation and asymptotically equal interpolations in model-based geostatistics. *Journal of the American Statistical Association 99*, 250–261.